arXiv:1604.00651v3 [quant-ph] 1 Nov 2016

A Direct Mapping of Max k-SAT and High Order Parity Checks to a Chimera Graph

N. Chancellor[†,1], S. Zohren[†,2,3], P. A. Warburton[4,5], S. C. Benjamin[2], S. Roberts[3]
[1] *Department of Physics, Durham University, South Road, Durham, UK*
[2] *Department of Materials, University of Oxford, Parks Road, Oxford, UK*
[3] *Department of Engineering Science, University of Oxford, Parks Road, Oxford, UK*
[4] *London Centre for Nanotechnology 19 Gordon St, London, UK*
[5] *Department of Electronic and Electrical Engineering, UCL, Torrington Place, London, UK*
[†] *Authors contributed equally. Correspondance to: nicholas.chancellor@gmail.com*
(Dated: November 2, 2016)

We demonstrate a direct mapping of max k-SAT problems (and weighted max k-SAT) to a Chimera graph, which is the non-planar hardware graph of the devices built by D-Wave Systems Inc. We further show that this mapping can be used to map a similar class of maximum satisfiability problems where the clauses are replaced by parity checks over potentially large numbers of bits. The latter is of specific interest for applications in decoding for communication. We discuss an example in which the decoding of a turbo code, which has been demonstrated to perform near the Shannon limit, can be mapped to a Chimera graph. The weighted max k-SAT problem is the most general class of satisfiability problems, so our result effectively demonstrates how any satisfiability problem may be directly mapped to a Chimera graph. Our methods faithfully reproduce the low energy spectrum of the target problems, so therefore may also be used for maximum entropy inference.

PACS numbers:

Introduction

Many interesting computer science problems have been shown to be directly mappable to finding the ground state of a Ising spin model on the hardware graph of the devices by D-Wave systems Inc. [1], the Chimera graph (see Figure 1). Examples include Maximum-Weight Independent Set, Exact Cover, and 3-SAT Problems [2]. Technically, by virtue of being NP-complete, an efficient mapping of any one of these problems implies that any other NP-complete problem can also be mapped to finding the ground state of a spin model on the Chimera graph. In practice however, such an indirect mapping is likely to be impractical given the current limitations of size and energy scales on real devices. It is for this reason that we are interested in direct mappings of interesting problems onto the Chimera graph, and why a direct mapping of a very general problem such as max k-SAT is of interest.

The ultimate reason why we are interested in mapping to the Chimera graph is to more efficiently map the 2-body Ising models which implement these problems to the hardware graph of the D-Wave annealers, although modifications to the ideas given here may be useful for mapping to other graphs which are made of tilings of locally non-planar graphs. Much of the attention in developing minor embeddings for the Chimera graph has been on increasing connectivity. In particular, a fully connected graph can always be mapped to a Chimera graph using minor embedding [11]. Recently, this focus has also included purpose-built architectures different from the Chimera graph [12] (see also [13] and [14] for related work). Here our focus is on higher order terms which are necessary to implement clauses and parity terms which involve more than two variables. However, as we will see later we will be using some of the minor embedding techniques for fully connected graphs to construct the embeddings of such terms.

There have been many promising advances in quantum annealing, since the idea that quantum fluctuations could help explore rough energy landscapes [15], through the algorithm first being explicitly proposed [16], further refined [17], and the basic concepts demonstrated experimentally in a condensed matter system [18]. Recently both entanglement [19] and tunneling [20, 22] have been experimentally demonstrated on programmable annealing processors. Given these encouraging results, it is desirable to propose new problem types for these machines to solve. This is interesting both from the viewpoint of possible eventual commercial applications, as well as providing access to new problems sets for benchmarking. For an overview of some aspects of quantum annealing, please see [21].

Another method of mapping problems onto the Chimera graph, is the one employed by [3]. This method uses numerical algorithms, often heuristic ones in practice, to map the problems. For reasons which we will discuss later, the low density parity check code (LDPC) decoding done in [3] can be thought of as a mapping of a weighted SAT problem into the Chimera graph. One crucial difference however, is that this LDPC decoding *cannot* be thought of as mapping a max-SAT problem, as our method does. For other examples of problem mappings see [23, 24].

Also unlike this mapping, our technique can not only be used for optimization tasks but also for sampling. This is important considering that recently there has been much interest in using D-wave for sampling applications, especially in the context of training Boltzmann machines [4–6], but also in the context of message decoding [7]. There are also many other examples in which maximum

entropy inference, which relies on sampling approximate thermal distributions can be applied in fields as varied as finance [8], ecology [9], and computational linguistics [10]. Even more powerful problem embeddings can probably be created by combining the ideas presented here with the powerful numerical techniques used in [3] .

We should also compare our paper to other recent work on using quantum annealing to solve satisfiability problems [25]. This work examines the use of quantum annealers to build SAT filters, which require a relatively large number of disparate solutions of a SAT problem to construct. This work finds that quantum annealing is not a suitable method compared to classical methods. Building a SAT solver is significantly different than solving such a problem directly, as we discuss in this paper, and is most appropriately classified as a variant of #SAT. Furthermore, SAT filters are not directly applicable to max-SAT problems as we study here, and so should be regarded as related, but very much distinct. It is worth remarking that the methods we give here could still potentially be useful for #SAT type problems such as SAT filter construction if the annealer were used to perform a hybrid algorithm as suggested in [26–28] rather than the standard quantum annealing algorithm.

For a review on boolean satisfiability, we point the reader to [29], and for max-SAT in particular to [30].

Implementing Clauses

Basic Operations

Any Boolean clause can always be written out as logical AND operations performed on strings of logical OR operators performed on bit values or the logical negation of bit values, e.g. $(a_1 \,\mathrm{OR}\, a_2 ...)\mathrm{AND}(\mathrm{NOT}\, a_1 \,\mathrm{OR}\, a_5 ...)$. In the following we denote AND by $\wedge$, OR by $\vee$ and negation by $\neg$. A general clause is thus of the form

$$(a_1^{(1)} \vee a_2^{(1)} ...) \wedge (a_1^{(2)} \vee a_2^{(2)} ...) \wedge ..., \tag{1}$$

where $a_i^{(l)} \in \{a_1, a_2, a_2, ...\} \cup \{\neg a_1, \neg a_2, \neg a_3, ...\}$. All that is needed to implement arbitrary Boolean clauses is therefore the ability to implement clauses of the form $a^{(l)} := (a_1^{(l)} \vee a_2^{(l)} ...)$. To implement a SAT problem in terms of energy computation, we could construct such a term by enforcing a penalty of the form,

$$Pen(\{a^{(l)}\}) \begin{cases} \geq g & a_i^{(l)} = 0,\ \forall i \\ = 0 & \text{otherwise.} \end{cases} \tag{2}$$

One can then construct a SAT problem by summing many such penalties and obtaining an energy $E = \sum_l Pen(\{a^{(l)}\})$. If one or more bit-strings exist where $E = 0$ then a set of clauses is satisfiable, but otherwise it is not. In the case where the clauses are satisfiable, the bit-strings which yield $E = 0$ are the ones which satisfy the clauses. However, because the penalties are unknown and are set to simply be an arbitrary value greater than or equal to g, the energies of states with $E > 0$ are meaningless. If no bit-string can satisfy all clauses, the lowest E state is *not* necessarily the one which satisfies the most clauses, and this is therefore *not* a valid construction of a max-SAT problem.

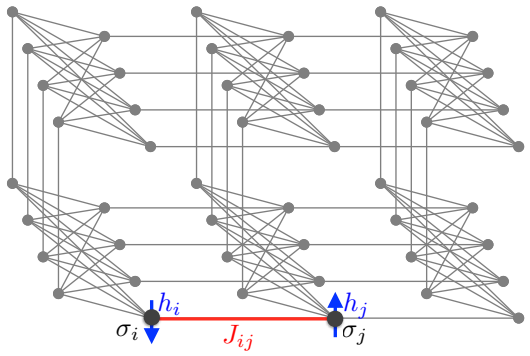

FIG. 1: A region of a Chimera graph containing 3×2 unit cells. Each vertex corresponds to a spin variable. One can adjust the magnetic fields h_i at each spin as well as the pairwise couplings J_{ij} between spins which are adjacent to each other in the graph.

However, if we consider terms which give all violated clauses the same energy penalty,

$$Spec(\{a^{(l)}\}) = \begin{cases} g & a_i^{(l)} = 0\ \forall i \\ 0 & \text{otherwise,} \end{cases} . \tag{3}$$

and similarly construct a total energy $E = \sum_l Spec(\{a^{(l)}\})$, then for $g > 0$ the minimum energy bit-string will always be the one which satisfies the most clauses, regardless of whether all clauses can be simultaneously satisfied. An energy penalty of this form therefore is a valid expression of a max-SAT problem. In this letter we show not only a natural way to express such penalties in terms of the Ising model, but also how such terms may be efficiently embedded into the D-Wave Chimera graph. It is worth pointing out that this can easily be even further generalized to a weighted version of the max k-SAT problem if a different value of g is chosen for different clauses $a^{(l)}$ in a controlled rather than arbitrary way.

To move from logical values to spin variables, we map each logical variable $a_i = 0$ to a spin variable with value $\sigma_i^z = -1$ and each logical variable $a_i = 1$ to a spin variable with value $\sigma_i^z = +1$. Negation of the logical variables is then implemented through gauges on the spin variables. More precisely, we map a_i to $c(i)\sigma_i^z$ with $c(i) = 1$ and $\neg a_i$ to $c(i)\sigma_i^z$ with $c(i) = -1$. Formally, we thus identify $a_i = \frac{1}{2}(1 + c(i)\sigma_i^z)$ with $c(i) = 1$ and $\neg a_i = \frac{1}{2}(1 + c(i)\sigma_i^z)$ with $c(i) = -1$.

We now consider how to implement a single SAT clause using an Ising spin Hamiltonian which might be implemented on an annealing machine like those of D-Wave

Logical bit values	Ancilla values	E
1111	0000	0
0111, 1011, 1101, 1110	0001	0
0011, 0101, 0110, 1001, 1010, 1100	0011	0
1000, 0100, 0010, 0001	0111	0
0000	1111	g

TABLE I: Summary of the possible configuration of bit-strings for the OR clause $a_1 \vee a_2 \vee a_3 \vee a_4$ together with the corresponding ancilla value and the energy (up to a constant offset).

systems Inc. Consider a Hamiltonian for spin variables of the form

$$\mathcal{H}^{(2)}_{\text{clause}} = J\sum_{i=1}^{k}\sum_{j=1}^{i-1} c(i)c(j)\sigma_i^z\sigma_j^z + h\sum_{i=1}^{k} c(i)\,\sigma_i^z + \\ +J^a\sum_{i=1}^{k}\sum_{j=1}^{k} c(i)\sigma_i^z\sigma_{j,a}^z + \sum_{i=1}^{k} h_i^a\sigma_{i,a}^z. \quad (4)$$

in which up to the gauge choice $c(i) \in \{-1, 1\}$, k *logical* spin variables σ_i^z are coupled to k ancilla spin variables – ancillae for short. The Hamiltonian above is similar to the Hamiltonian presented in [31] and in fact both are special cases of a more general construction presented below. To implement a single clause, we set $J = J^a$, $h = -J^a$, as well as $h_i^a = -J^a(2i - k) + q_i$ with

$$q_i = \begin{cases} g/2 & i = 1, \\ 0 & \text{otherwise}, \end{cases} \quad (5)$$

where $g/2 \ll J_a$. By the symmetry of this Hamiltonian, the effective energy penalty on the ancillae from the logical spin variables for being up or down will depend only on the total number of logical spin variables which are in agreement with the gauges, and not the specific arrangement. With this choice of q_i, all bit-strings will have the same energy unless $c(i)\sigma_i^z = -1, \forall i$ in which case the energy will be greater by g. Thus, up to an irrelevant total energy shift, (4) generates a single penalty term of the form of (3). Table I illustrates an example for four bits with all gauge values set to 1. Figure 2 shows the connectivity of the corresponding abstract spin graph. Shown are the four logical spin variables in green and the ancillae in red.

We have demonstrated above how to implement a single clause of the form $a^{(l)} = (a_1^{(l)} \vee a_2^{(l)} ...)$. Using this construction we can now implement clauses of the form $(a_1^{(1)} \vee a_2^{(1)} ...) \wedge (a_1^{(2)} \vee a_2^{(2)} ...) \wedge ...$ by superimposing the construction of the individual clauses $a^{(l)}$ on a common set of logical spin variables. To do so, we include the gauge variables in the coupling constants, i.e. associated with the clause $a^{(l)}$ the fields are $h_i^{(l)} = c^{(l)}(i)h^{(l)}$ and

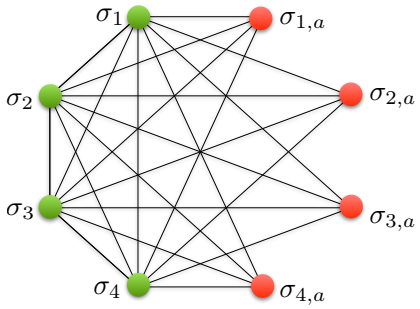

FIG. 2: Illustration of the spin graph corresponding to the implementation of the clause $\sigma_1 \vee \sigma_2 \vee \sigma_3 \vee \sigma_4$, as well as the parity checking clause $\sigma_1 \oplus \sigma_2 \oplus \sigma_3 \oplus \sigma_4$. Logical spin variables are shown as green vertices, ancillae as red vertices, non-zero couplings are shown as black edges and magnetic fields are not shown. Both clauses for OR and XOR only differ in the value of their fields.

the couplings are $J_{ij}^{(l)} = c^{(l)}(i)c^{(l)}(j)J^{(l)}$. The total fields applied to a spin variable σ_i^z are then the sum of the field contribution from each clause, i.e. $h_i = \sum_l h_i^{(l)}$ and similarly for the couplings $J_{ij} = \sum_l J_{ij}^{(l)}$. Note that the ancillae cannot be superimposed.

XOR Clauses and Parity Checks

We now describe another construction where instead of OR we have clauses constructed out of XOR relations. Such clauses are important in many message decoding applications. Furthermore, they provide an alternative method for implementing the above clauses. In particular, instead of implementing clauses of the form $a^{(j)} := (a_1^{(j)} \vee a_2^{(j)} ...)$ and superimposing those to construct the bigger problem $(a_1^{(1)} \vee a_2^{(1)} ...) \wedge (a_1^{(2)} \vee a_2^{(2)} ...) \wedge ...$, one can also implement bigger clauses directly. To do so, note that a logical AND operation can be expressed as a product of two such operators $a_i \wedge a_j = a_i\, a_j$ and the OR operation can be written as $a_i \vee a_j = a_i + a_j - 2\, a_i\, a_j$. Inserting the mapping to spin variables, $a_i \to \frac{1}{2}(\sigma_i^z + 1)$, we see that any Boolean clause can be rewritten in terms of spin variables, and the penalty Hamiltonians can alternatively be constructed using the methods given in [31] in terms of multi-body terms. A product of spins corresponds to a *parity term* or *parity checking clause*, which is represented by an XOR relation in the Boolean language, i.e. a term of the form $(a_1^{(l)} \oplus a_2^{(l)} \oplus a_3^{(l)} ...)$, where $\oplus$ denotes bitwise addition. Thus we have the mapping,

$$a_1^{(l)} \oplus a_2^{(l)} \oplus ... \oplus a_k^{(l)} \quad \leftrightarrow \quad c(1)\sigma_1^z c(2)\sigma_2^z ... c(k)\sigma_k^z \quad (6)$$

The spectrum $Spec(\{a_1^{(l)} \oplus ... \oplus a_k^{(l)}\})$ is mapped to the spectrum of the Hamiltonian

$$\mathcal{H}_k = \frac{g}{2} c(1)\sigma_1^z \; ... \; c(k)\sigma_k^z. \quad (7)$$

The abstract connectivity graph of this Hamiltonian is shown in Figure 2 for $k = 4$. We can use the Hamiltonian (4) to reproduce the (low-energy) spectrum of $\mathcal{H}_k$ where, as before $J = J^a$, $h = q_0 - J^a$, as well as $h_i^a = -J^a(2i - N) + q_i$, but q_i is instead chosen as

$$q_i = \begin{cases} q_0 + g/2 & N - i \text{ is odd}, \\ q_0 - g/2 & N - i \text{ is even}. \end{cases} \quad (8)$$

with $g/2 < q_0 \ll J_a$. This assignment of coupling constants is the same as used in [31].

It is also worth briefly pointing out that the mapping between logical and spin variables discussed above is completely invertible. Any problems which can be expressed as a sum over products of spin operators can therefore be written as a weighted sum over Boolean clauses plus a constant. Problems expressed as sums of parity checking clauses which are each weighted equally can therefore be regarded as a version of max k-SAT. In an upcoming work [34] we analyze the potential of using this way of implementing parity checking terms in terms of these weighted max-SAT implementations for inference in message decoding problems, in particular when applied to Low Density Parity Check (LDPC) codes and turbo codes. The advantage of our construction in comparison to an earlier construction of LDPC codes [3] is that it not only reproduces correctly the ground state but also the low-energy spectrum, thus permitting us to do sampling applications, such as the maximum entropy inference discussed in [7]. While we leave the extensive analysis of various message decoding problems to the upcoming work [34], later in this manuscript we give a brief description of how to implement turbo codes.

Other Clauses

We have already demonstrated that a simple clause can be expressed as an energy penalty by using k ancillae. This therefore allows a natural construction of any complex k bit clause by simply examining every possible bit string and penalizing it if it violates the clause. For large k this method is rather inefficient for arbitrary clauses, however, as the number of bit-strings which must be examined potentially grows as $k2^k$. For example, expressing a parity checking clauses using individual SAT clauses would require $k\,2^{\lfloor k/2 \rfloor}$ ancillae.

It is therefore worthwhile to briefly address how one might go about constructing methods for implementing clauses which cannot be easily expressed using the methods previously discussed more efficiently. Firstly we note that the previous constructions can be generalized by choosing q_i in (5) and (8) differently. By doing this we can implement any clause which is symmetric under permutation of any of the bits. Such a clause will be defined by a vector $f_i \in \{0, g/2\}$ with $i = 0, ..., k$ which is 0 if the set of bit-strings with i bits equal to 1 satisfies the clause and $g/2$ otherwise. We now define

$$q_i = q_0 + f_{i-1} - f_i, \quad i = 1, ..., k \quad (9)$$

with $g/2 < q_0 \ll J^a$. We further observe that the gauges $c(i) \in \{-1, 1\}$ in (4) allow us to define such clauses which are symmetric *in any gauge*, and potentially to combine more than one of this type of clause constructed in multiple different gauges. These more complicated constructions should allow many different clauses involving relatively large numbers of bits to be implemented more efficiently than the method given earlier. In practice one would probably want to construct numerical algorithms to find more optimal implementations of arbitrary high k clauses, but this is beyond the scope of the current letter.

Special Cases

It is worth briefly mentioning a couple of special cases, in which clauses can be expressed more efficiently than the ways discussed earlier in this section. Because the fields and couplers already act as one and two bit parity checking clauses (a field, which gives a different energy for a 1 and 0 state is nothing more than a single bit parity check), all that is needed to construct an arbitrary three bit clause is a three bit parity checking clause. Using the construction given previously for such clauses requires three ancillae. This can however be reduced to a single ancilla by choosing $h = g$, $J_a = 2J > |h|$, and $h_a = 2h$. *Any* 3 bit clause can therefore be constructed using only a single ancilla. The corresponding abstract spin graph is shown in Fig. 3 (a).

Furthermore, clauses of the form $(a_1^{(l)} \wedge ... \wedge a_k^{(l)})$ can be expressed using only a single ancilla per sub-clause, regardless of k. Consider the simpler Hamiltonian of the form

$$\mathcal{H}^{(2)}_{\text{simple}} = h \sum_{i=1}^{k} c(i)\, \sigma_i^z + J^a \sum_{i=1}^{k} c(i) \sigma_i^z \sigma_a^z + h^a \sigma_a^z \quad (10)$$

in which up to the gauge choice $c(i) \in \{-1, 1\}$, k logical spin variables σ_i^z are coupled with equal strength J^a to the same ancilla spin variable σ_a^z. The connectivity of the corresponding spin graph is shown in Fig. 3 (b) for the case $k = 4$. We choose h^a so that if all of the logical bits σ_i^z match $c(i)$, the ancilla bit will be down. This can be achieved as follows: Consider choosing $h_a = J^a\, k + q$. The ground state of the ancilla will be the upward orientation unless all logical bits cooperate to counteract the field in which case it will be downward. We can further choose the couplers between the logical bits such that the energy of the ancilla-logical couplers exactly cancels the energy from the couplers between the logical qubits. We now further set $J^a = -h$ and $q = -g/2$, leading to

$$h_a = -h\, k - g/2, \quad J^a = -h \quad (11)$$

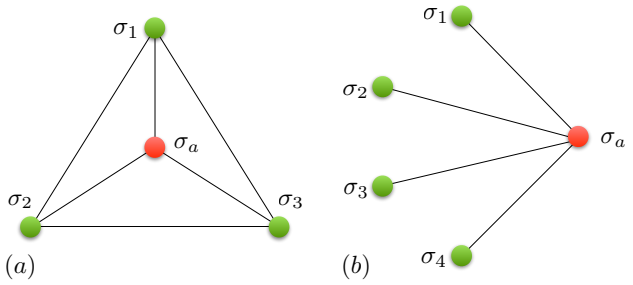

FIG. 3: (a) Illustration of the abstract spin graph corresponding to the implementation of the clause $\sigma_1 \oplus \sigma_2 \oplus \sigma_3$ using only a single ancilla. (b) Illustration of the abstract spin graph for the clause $\sigma_1 \wedge \sigma_2 \wedge \sigma_3 \wedge \sigma_4$.

which (up to a constant offset) yields a total energy of zero iff $c(i)\sigma_i^z = +1, \forall i$ and g otherwise, which is exactly the spectrum of a clause of the form $(a_1^{(l)} \wedge ... \wedge a_k^{(l)})$.

Embedding in the Chimera Graph

Let us start by considering how to embed both simple clauses and parity checks represented by the connectivity graph shown in Fig. 2 into a patch of a Chimera graph. This embedding is shown in Fig. 4. The abstract graph shown in Fig. 2 is already reduced to have only standard two-body interactions between neighboring spins as is the case for the Chimera graph. The only difficulty in embedding this abstract graph is the fact that its connectivity is higher than that of the Chimera graph. In particular, the embedding of the clauses of order k involves a fully connected graph of the k logical spin variables. Higher connectivity can be achieved at the price of an overhead in the number of spin variables by 'identifying' different spin variables through a very strong link. In other words, two spins are coupled through a ferromagnetic link of strength $|J^\infty|$ which is much larger than all the other couplings, ensuring that both spins always have the same value. This can be seen in the embedding of the clauses in Fig. 4, where the strong links identifying logical spin variables are shown as thick green edges.

To generalize the above embedding for a larger number of variables and clauses one can employ the minor embedding of a fully connected graph which was introduced in [11]. Since clauses can be superimposed, it suffices to have a single minor embedding of all logical spin variables. While one could also include all ancillae in a single fully connected graph and then set unused edges to zero, this would not be very efficient. A more efficient way to do this is to extend each of the logical spin variables as a string of physical spin variables coming out of one side of the fully connected embedding of all logical spin variables, with each of the ancillae as an embedding chain crossing all of them. This is illustrated in Fig. 5. In the lower three rows of unit cells we see a minor embedding of a fully connected graph between 12 logical spin variables. In the upper part of the figure, each of the 12 spin variables has outgoing chains of 'identified' spins which are then connected to the ancillae. This plot illustrates a specific example of an embedding to perform some of the parity checks for the turbo code example given previously which is explained in detail in that section.

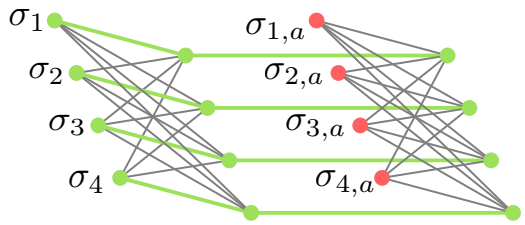

FIG. 4: Illustration of the spin graph corresponding to the minor embedding of the clause $\sigma_1 \vee \sigma_2 \vee \sigma_3 \vee \sigma_4$ as well as the parity checking clause $\sigma_1 \oplus \sigma_2 \oplus \sigma_3 \oplus \sigma_4$ in the Chimera graph. The unit cell on the left implements the fully connected graph amongst the for logical variables, while the unit cell on the right implements the ancillae. Strong coupling edges between vertices (shown as thick green edges) are used to "identify" spin variables.

The method proposed in the previous paragraph works well for embedding any number of arbitrary overlapping subgraphs of the form shown in Fig. 2. For problems requiring a large number of ancillae however, embeddings with this method will only occupy a long relatively thin strip of the Chimera graph. In practice, real devices tend to be designed with an aspect ratio close to 1:1, which maximizes the tree-width for given number of qubits. Because of this, the way to embed into a real device will be to use a serpentine pattern, which uses fully connected graphs as 'corners' at the end of each row and allow the embedding to efficiently fill the graph.

Scaling

Let us now consider the scaling of the total number of physical qubits, N_{phys} required to embed an instance of a max k-SAT problem using penalty terms of the form (3) only. We shall first assume that the long stretches of linear embedding chains which cross the ancillae will dominate. In this case the total number of physical qubits will scale as,

$$N_{phys} \propto N_{log} \langle k \rangle \, c \tag{12}$$

where N_{log} is the number of logical qubits, c is the number of clauses, and $\langle k \rangle$ is the mean number of bits per clause. It has been demonstrated, for instance with max 2-SAT [32, 33] that the typical hardness of a problem is determined by the ratio $r = c/N_{log}$. If clauses are too sparse, then it will typically be easy to satisfy them all simultaneously, however if clauses are too dense, the problem again becomes easy because no solution will be able to satisfy very many of the clauses and almost any

random bit string will be a good solution. Based on this reasoning, the value of r which gives the hardest typical problems should not vary too much from $r \approx O(1)$ and therefore if we want scaling for the hardest problems r can be treated as roughly independent of $\langle k \rangle$, the mean number of bits in a clause.

If we are interested in using annealers to solve problems drawn from a typically hard set of problems, the number of physical qubits required should scale roughly as,

$$N_{phys} \propto N_{log}^2 \ \langle k \rangle \ r \tag{13}$$

which, is the same scaling as the minor embedding for a fully connected graph proposed in [11].

This scaling only applies for those clauses which can be expressed as terms of the form (3). In general one can imagine much more complicated clauses. Using the more naive method of constructing a clause piecewise by penalizing each bitstring on k bits can at most require $2^k - 1$ subclauses to implement. Each of these subclauses will require an ancilla to implement. For this reason the average number of ancillae to implement naively is $\langle N_{anc}(k) \rangle \leq 2^k - 1$. The actual value of $\langle N_{anc}(k) \rangle$ is likely to be highly dependent on both the specific max-SAT problem and the algorithm to implement the clauses as discussed previously.

For arbitrary clauses, involving a maximum of k bits, the worst case scenario scaling for typically hard problems is therefore,

$$N_{phys} \propto N_{log}^2 \, r \ \langle N_{anc}(k) \rangle \, , \tag{14}$$

which scales exponentially with the clause length k, but still only scales quadratically with N_{log}, given a fixed value of k. This scaling still assumes we are in the typically hard regime. However, we can further calculate the absolutely worst case, which one clause needs to be constructed on every subset of k or fewer bits. Assuming that $c \gg k$, then the leading order scaling in the number of ancilla required will be N_{log}^k, and the overall scaling will be, assuming $k \geq 3$,

$$N_{phys} \propto N_{log}^{k+1}, \tag{15}$$

which is still polynomial in N_{log} for fixed k, however the rapidly growing power will place limits on practical realizations of such pathological problems. It worth pointing out here that while such problems can be mathematically constructed, it is not clear that any problems whose embedding scales like (15) are actually of any practical interest, or indeed even that hard problems which scale like this exist.

It is worth noting that while it is possible to implement 3 bit clauses with only a single ancilla, this technique does not have a meaningful effect on the overall scaling.

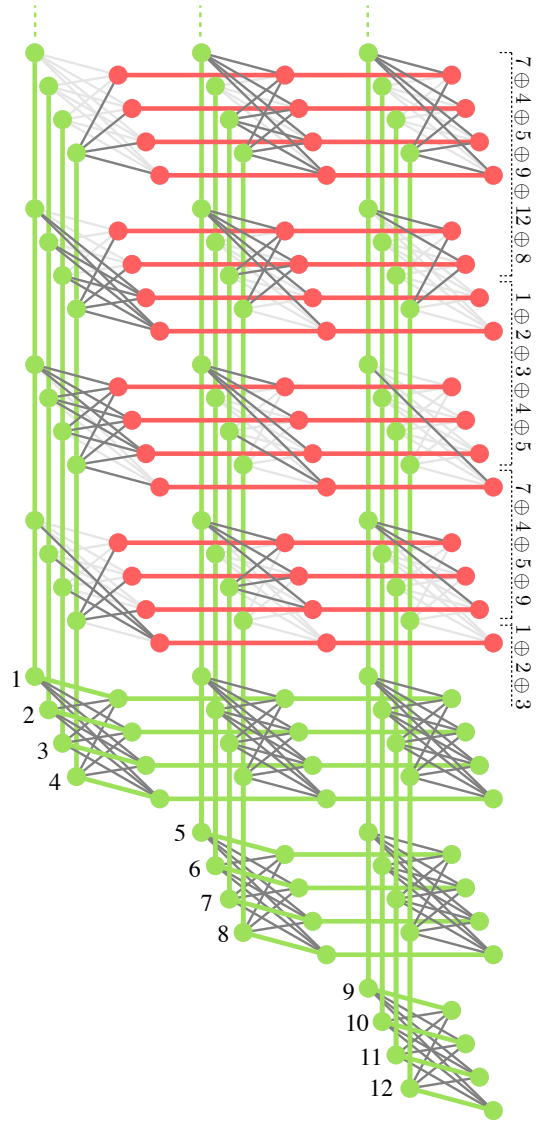

FIG. 5: Embedding of a few clauses of a 12 bit problem in a Chimera graph. The actual problem is part of a turbo code involving 12 message bits and 12 parity bits (of which eight are omitted in the figure). The bottom 3 rows implement a fully connected graph amongst the 12 message variables, which allows for arbitrary 2 bit clauses. The upper 4 rows illustrate how the ancillae can be embedded. To implement a clause using our method, the number of ancillae needs to be equal to the number of logical bits in the clause and each row of Chimera unit cells provides 4 ancillae.

Application to Turbo Codes

As a practical example of scaling, let us consider decoding a turbo code. We restrict ourself to the basic construction and leave a detailed analysis to a forthcoming work [34]. While this problem is most naturally stated in terms of Ising spins, we have already shown that this problem maps to a max k-SAT. Turbo codes are a class of so-called convolutional code which have many real life ap-

plications in communications due to the ability to achieve near Shannon limit performance [35–37]. The Hamiltonian for decoding of a general turbo code can be written in the following way,

$$H(\rho) = f\sum_{i=1}^{K}\rho_i\sigma_i^z + f\sum_{i=1}^{K}\rho_{i+K}\left(\mathbb{I}(i \text{ odd})\prod_{j=1}^{i}\sigma_j^z + \right. \\ \left. +\mathbb{I}(i \text{ even})\prod_{j=1}^{i}\sigma_{p(j)}^z\right) \quad (16)$$

where $p(j)$ is a random permutation, and $\rho_i \in \{-1, 1\}$, with $i = 1, ..., 2K$ the list of received values of the message and the corresponding parity checks. Furthermore, f is related to the noise model; in the case of a binary symmetric channel which corrupts a value with probability p, one has

$$f = -\frac{1}{2}\log\left(\frac{1-p}{p}\right). \quad (17)$$

The values ρ_i, with $i = 1, ..., K$ represent the received values corresponding to the original message, while ρ_i, with $i = K+1, ..., 2K$ are the received values corresponding to the parity checks. The parity checks are performed in a nested structure, firstly one transmits parity checks on odd numbers of variables with respect to the original ordering of the spin variables, i.e. σ_1^z, $\sigma_1^z\sigma_2^z\sigma_3^z$, $\sigma_1^z\sigma_2^z\sigma_3^z\sigma_4^z\sigma_5^z$, etc. and secondly one transmits parity checks of even number of variables on the permuted variables, i.e. $\sigma_{p(1)}^z\sigma_{p(2)}^z$, $\sigma_{p(1)}^z\sigma_{p(2)}^z\sigma_{p(3)}^z\sigma_{p(4)}^z$, etc.

The number of ancillae required for each parity checking clause scales like the length of the clause, and there is one clause for every possible clause length in (16). The number of ancilla required to decode the turbo code therefore scales as $N_{anc} \propto N_{log}^2$. The total number of physical bits will therefore scale as $N_{phys} \propto N_{log}^3$ for turbo code decoding.

To give an explicit example, Fig. 5 shows part of the embedding of a turbo code. We have $K = 12$ message variables and an equal number of parity checks. The message variables are represented by the logical spin variables (green) which are fully connected. On the outgoing chains we couple the ancillae corresponding to the parity checks. Note that parity checks of order 1 and 2 can be directly implemented on the fully connected graph. The next highest order is the parity check of order three on the variables 1, 2 and 3 which is implemented using the efficient embedding using a single ancilla. Next we have a parity check on the four variables $p(1), p(2), p(3), p(4)$, where for concreteness we chose $p(1) = 7, p(2) = 4, p(3) = 5, p(4) = 9$. The figure also shows the next two parity checks of order 5, involving variables $1, 2, 3, 4, 5$, and of order 6, involving $p(1), p(2), p(3), p(4), p(5), p(6)$, where in the example we chose $p(5) = 12, p(6) = 8$.

Conclusion

We have demonstrated a method to embed max k-SAT (and weighted max k-SAT), including the problem of finding a bit-string which satisfies a maximum number of parity checks, which can be expressed in terms of clauses involving XOR. We further demonstrate a very efficient way to implement such clauses, which have important applications in communications, we discuss the specific example of turbo code decoding.

The weighted max k-SAT problem is the most general satisfiability problem, so we have therefore demonstrated how *any* satisfiability problem can be directly mapped into a Chimera graph. One particularly interesting application of this is parity checking, which could lead to important applications in communications. Furthermore, the methods given here reproduce the low energy spectrum of the problem, with energies corresponding to the number of clauses which are unsatisfied. This means that as well as finding the lowest energy solution, these techniques are compatible with maximum entropy inference applications.

Our method gives a direct construction of the problem Hamiltonians without the need for numerically expensive classical calculations. However the methods given here could probably be made more powerful if integrated into the already powerful and growing numerical toolset which is currently used to map problems into a Chimera for real calculations, for example in [3]. The embedding illustrated here was chosen for its generality, but for specific problems it is unlikely that every ancilla will have to couple to every logical qubit, or that 2-qubit couplers between every logical qubit will be necessary. Therefore, embedding efficiency gains are likely to be possible through numerical optimization. Furthermore it would be interesting to explore how these methods can be generalized to more sparse graphs and more general clause types.

Acknowledgments

N.C. was supported by EPSRC (grant ref: EP/L022303/1). S.Z. acknowledges support by Nokia Technologies, Lockheed Martin and the University of Oxford through the Quantum Optimisation and Machine Learning (QuOpaL) Project. P.W. was supported by Lockheed Martin and by EPSRC (grant refs: EP/K004506/1 and EP/H005544/1). SCB is supported by the EPSRC National Quantum Technology Hub in Networked Quantum Information Technologies (grant ref: EP/M013243/1).

Author Contributions

SZ, NC and PW came up with the initial concept and wrote the paper, with SZ creating the figures. SB and

SR helped to supervise and were involved in developing the scientific content and editing the manuscript along with all other authors.

Competing Interests

The authors declare no competing interests.

[1] http://www.dwavesys.com/ accessed: March 21, 2016

[2] Choi, V. Adiabatic Quantum Algorithms for the NP-Complete Maximum-Weight Independent Set, Exact Cover and 3SAT Problems arXiv:1004.2226 (2004).

[3] Bian Z. et. al. Discrete optimization using Quantum Annealing on sparse Ising models *Frontiers in Physics* **2** 00056, (2014).

[4] Adachi, S. H. and Henderson, M. P., Application of Quantum Annealing to Training of Deep Neural Networks, arXiv:1510.06356 (2015).

[5] Benedetti, M., Realpe-Gmez, J., Biswas, R. and Perdomo-Ortiz, A. Estimation of effective temperatures in quantum annealers for sampling applications: A case study towards deep learning arXiv:1510.07611 (2015).

[6] Amin, M. H., Andriyash E., Rolfe, J., Kulchytskyy B. and Melko, R. Quantum Boltzmann Machine, arXiv:1601.02036 (2016).

[7] Chancellor, N., Szoke, S., Vinci, W., Aeppli, G. and Warburton, P. A. Maximum-Entropy Inference with a Programmable Annealer *Scientific Reports* **6** 22318 (2016).

[8] Mistrulli, P. E. Assessing financial contagion in the interbank market: Maximum entropy versus observed interbank lending patterns. *Journal of Banking and Finance* **35** 1114-1127 (2011).

[9] Phillips, S. J. et. al. Maximum entropy modeling of species geographic distributions. *Ecological Modelling* **190** 231-259 (2006).

[10] Berger, A. L. et al. A maximum entropy approach to natural language processing. *Computational Linguistics 22* DOI:10.1.1.103.7637 (1996).

[11] Choi, V. Minor-embedding in adiabatic quantum computation: II. Minor-universal graph design *Quantum Inf Process* **10** 343353 (2011).

[12] Lechner, W., Huke, P. and Zoller, P., A quantum annealing architecture with all-to-all connectivity from local interactions *Sci. Adv.* **1** e1500838 (2015).

[13] Albash T., Vinci W., and Lidar D. A., Simulated Quantum Annealing with Two All-to-All Connectivity Schemes, arXiv:1603.03755 (2016).

[14] Rocchetto, A., Benjamin, S. C. and Li, Y., Stabiliser sets and physical ancillas for the Lechner-Hauke-Zoller Annealer, arXiv:1603.08554 (2016).

[15] P. Ray, B. K. Chakrabarti and A. Chakrabarti, Sherrington-Kirkpatrick model in a transverse field: Absence of replica symmetry breaking due to quantum fluctuations *Phys. Rev. B* **39** 11828 (1989)

[16] Finilla A. B., Gomez M.A., Sebenik C, Doll D. J. Quantum annealing: a new method for minimizing multidimensional functions. *Chem Phys Lett.* **219** 343348 (1994).

[17] Kadowaki, T. Nishimori, H. Quantum annealing in the transverse Ising model *Phys. Rev. E* **58** 5355(1998).

[18] Brooke, J., Bitko, D., Rosenbaum, T. F. and Aeppli, G. Quantum annealing of a disordered magnet. *Science* **284** 779781 (1999).

[19] Lanting, T. et al. Entanglement in a quantum annealing processor *Phys. Rev. X* **4** 021041 (2014).

[20] Boixo, S. et. al. Computational Role of Multiqubit Tunneling in a Quantum Annealer *Nature Communications* **7** 10327 (2016).

[21] Das, A. Chakrabarti, B. K. Colloquium: Quantum annealing and analog quantum computation *Rev. Mod. Phys.* **80**, 1061 (2008).

[22] Boixo S. et. al. Computational Role of Collective Tunneling in a Quantum Annealer arXiv:1411.4036 (2014).

[23] J. D. Biamonte Nonperturbative k-body to two-body commuting conversion Hamiltonians and embedding problem instances into Ising spins *Phys. Rev. A* **77** 052331 (2008).

[24] J. D. Whitfield, M. Faccin, and J. D. Biamonte Ground-state spin logic *EPL* **99** 57004 (2012)

[25] Azinovic, M. Herr, D. Heim, B. Brown, E. Troyer, M. Assessment of Quantum Annealing for the Construction of Satisfiability Filters, arXiv:1607.03329 (2016).

[26] Neven H. et. al. Opening remarks: Why we believe quantum annealing will succeed, presented at: AQC 2016 (2016).

[27] Chancellor N. Modernizing quantum annealing using local searches. arXiv:quant-ph/1606.06833, (2016).

[28] Chancellor N. Modernizing Quantum Annealing II: Genetic Algorithms and Inference arXiv:1609.05875 (2016).

[29] Malik, S. Zhang, L. Boolean Satisfiability: From Theoretical Hardness to Practical Success *CACM* **52** 8 (2009).

[30] Stutzle, T., Hoos, H., Roli, A. A review of the literature on local search algorithms for MAX-SAT. *Rapport technique AIDA-01-02, Intellectics Group, Darmstadt University of Technology, Germany* (2001).

[31] Chancellor, N., Zohren, S., Warburton, P. A., Circuit design for multi-body interactions in superconducting quantum annealing system with applications to a scalable architecture, arXiv:1603.09521 (2016).

[32] Coppersmith, D., Gamarnik, D. Hajiaghayi, M. T. and Sorkin, G. B. Random MAX SAT, random MAX CUT, and their phase transitions *Random Struct. Alg.* **24** 502545. (2004).

[33] Santra, S., Quiroz, G., Ver Steeg, G. and Lidar, D. MAX 2-SAT with up to 108 qubits *New J. Phys.* **16** 045006 (2014).

[34] Chancellor, N., Zohren, S., Warburton, P. A., Benjamin, S. C. and Roberts, S., Inference problems with higher locality on a quantum annealing device: From Hamming codes to LDPC and turbo codes, to appear (2016).

[35] Berrou, C., Glavieux, A., Near optimum error correct-

ing coding and decoding: Turbo-codes, *IEEE Trans. on Comm.* **44** 12611271 (1996).

[36] Berrou, C., Glavieux, A., Thitimajshima, P., Near Shannon limit error-correcting coding and decoding: Turbo-codes, *Proc. 1993 IEEE Int. Conf. on Communications, Geneva, Switzerland* 1064, (1993).

[37] MacKay, D. *Information Theory, Inference and Learning Algorithms*, Cambridge University Press, (2003).